arXiv:1601.06652v1 [cs.SD] 25 Jan 2016

A PERCEPTUALLY MOTIVATED FILTER BANK WITH PERFECT RECONSTRUCTION FOR AUDIO SIGNAL PROCESSING

THIBAUD NECCIARI, NICKI HOLIGHAUS, PETER BALAZS AND ZDENĚK PRŮŠA

Abstract. Many audio applications rely on filter banks (FBs) to analyze, process, and re-synthesize sounds. To approximate the auditory frequency resolution in the signal chain, some applications rely on perceptually motivated FBs, the gammatone FB being a popular example. However, most perceptually motivated FBs only allow partial signal reconstruction at high redundancies and/or do not have good resistance to sub-channel processing. This paper introduces an oversampled perceptually motivated FB enabling perfect reconstruction, efficient FB design, and adaptable redundancy. The filters are directly constructed in the frequency domain and linearly distributed on a perceptual frequency scale (e.g. ERB, Bark, or Mel scale). The proposed design allows for various filter shapes, uniform or non-uniform FB setting, and large down-sampling factors. For redundancies ≥ 3 perfect reconstruction is achieved by computing the canonical dual FB analytically. For lower redundancies perfect reconstruction is achieved using an iterative method. Experiments show performance improvements of the proposed approach when compared to the gammatone FB in terms of reconstruction error and resistance to sub-channel processing, especially at low redundancies.

1. Introduction

Time-frequency (TF) transforms like the short-time Fourier or wavelets transforms play a major role in audio signal processing. They allow decomposing any signal into a set of elementary functions with good TF localization and achieving perfect reconstruction if the transform parameters are chosen appropriately (e.g. [15]). Therefore, they constitute ideal tools to analyze, process and re-synthesize sounds. Accordingly, applications like audio coding [6, 47], audio transformations [40, 45], sparsity [5, 38], source separation [16, 26], speech processing [20, 32], de-noising [12, 31], or optimization of acoustical measurements [30], among others, rely on TF decompositions to perform sub-channel processing and reconstruct the signal from the modified TF components. In such applications, TF transforms are usually implemented as filter banks (FBs) where the set of analysis filters defines the elementary functions and the set of synthesis filters allows for signal reconstruction. The TF concentration of the filters together with the downsampling factors in the sub-bands define the TF resolution and redundancy of the transform. FBs come in various flavors and have been extensively treated in the literature (e.g. [1, 50]). Note that the mathematical theory of frames constitutes an interesting alternative background for the interpretation and implementation of FBs (see e.g. [2,3,8,10,14] and Appendix A).

Because sub-channel processing may introduce audible distortions in the reconstructed signal, particularly if the sub-bands are not equally processed, important

requirements of audio applications include: a *strong* stability (i.e. the coefficients are bounded if and only if the signal is bounded, i.e. the FB and its inverse are BIBO-stable), perfect reconstruction property of the analysis-synthesis system (i.e. when no sub-channel processing is performed)[1], resistance to noise, and adequate aliasing suppression in each sub-band. To limit the computational costs, certain applications also require low redundancy, that is a small number of sub-bands with large downsampling factors. While sub-sampling the sub-bands is usually not required in speech processing where signals are often short and sampled at low rates (typically 8 kHz), sub-sampling is of high interest to music processing where signals are few seconds long and sampled at high rates ($\geq$ 44.1 kHz).

Although many applications still use transforms with fixed resolution (e.g. short-time Fourier or modified cosine transforms), there is a strong desire in audio processing to analyze sounds in a manner similar to that of the human ear. Since the auditory TF resolution varies with frequency, this implies using a transform with variable resolution. This purpose has lead to the design of so-called auditory FBs (e.g. [22, 24, 28]) or perceptually motivated FBs (e.g. [9, 43, 52]). Perceptually motivated FBs are usually intended as signal processing tools and, as such, they are linear, partially invertible, and have good aliasing suppression but only approximate the auditory frequency resolution. In contrast, auditory FBs are usually intended as perceptual analysis tools and, as such, they attempt to reproduce the nonlinear processing in the auditory system, to the detriment of perfect reconstruction, aliasing suppression, and computational efficiency. A *linear* and partially invertible auditory FB became popular in audio signal processing, though, namely the gammatone FB (e.g. [22, 39]). Gammatone filters approximate well the auditory TF resolution at low to moderate sound levels and are easy to implement as FIR or IIR filters [22, 27, 29]. A wide range of audio applications thus implements gammatone FBs, for instance source separation [16], speech processing [20, 42, 54], or music information retrieval [51]. Still, gammatone FBs do not satisfy all requirements of audio applications as they neither provide perfect reconstruction nor good aliasing suppression (see Sec. 2.3).

To fulfill the requirements of audio applications, this paper introduces an oversampled perceptually motivated FB enabling perfect reconstruction, efficient FB design, and adaptable resolution and redundancy. In the proposed approach, the filters are directly constructed in the frequency domain and linearly distributed on a perceptual frequency scale (e.g. ERB, Bark, or Mel scale). The proposed design allows for various filter shapes, uniform or non-uniform FB setting, and large downsampling factors. For redundancies ≥ 3 perfect reconstruction is achieved by computing the dual FB directly. For lower redundancies (down to $\approx$1.1) perfect reconstruction is achieved using an iterative method. Experiments show the better performance of the proposed approach with respect to the commonly-used gammatone FB in terms of reconstruction error and resistance to sub-channel processing, especially at low redundancies.

The paper is organized as follows. The next section briefly describes the properties of auditory frequency selectivity and perceptual frequency scales, and reviews recent works related to the present study. Section 3 introduces the analytical and implementation properties of the proposed "AUDlet" FB. Finally, simulations are

[1]Note that because of the quantization process, the perfect reconstruction requirement might be violated for lossy coding applications.

performed in Section 4 to show the performance of the AUDlet FB in terms of signal representation, reconstruction error, and as an audio processing tool. In the appendix, general results on non-uniform FBs and their connection to the theory of frames are recalled for the better understanding of the AUDlet FB construction.

2. Background

2.1. **Auditory Frequency Selectivity.** The frequency selectivity of the auditory system can be modeled in a first approximation as a bank of bandpass filters, named "critical bands" or "auditory filters", that are related to the frequency-to-place transformation in the cochlea (see e.g. [33, Chap. 3] for a review). Briefly, when a sound reaches the ear it produces a vibration pattern on the basilar membrane. The position and width of this pattern along the membrane depend on the spectral content of the sound. Accordingly, the center frequency and bandwidth of the auditory filters respectively approximate the place and width of excitation on the basilar membrane. Noteworthy, the width of excitation depends on level as well: patterns become wider and asymmetric as sound level increases (e.g. [18]). Several auditory filter models have been proposed based on the results from masking experiments [29]. A popular auditory filter model is the gammatone filter [39]. Although gammatone filters do not capture the level dependency of the actual auditory filters, their ease of implementation in the time or Laplace domains made them popular in audio signal processing (e.g. [16,20,51,54]). More realistic auditory filter models are, for instance, the roex and gammachirp filters [18, 49].

2.2. **Perceptual Frequency Scales.** To analyze sounds using a frequency resolution that mimics that of the ear or to match the spectral content of a sound to an auditory sensation (e.g. pitch or loudness), a mapping between the linear frequency domain and the nonlinear perceptual domain is required. This mapping is provided by perceptual frequency scales developed based on psychoacoustics experiments. We mention below three scales that are commonly used in hearing science and audio signal processing, namely the Bark, ERB, and Mel scales. To describe the different mappings we introduce the function $F : \xi \to \mathrm{AUD}$ where ξ is frequency in Hz and AUD is an auditory unit that depends on the scale.

2.2.1. *The Bark Scale.* Directly originates from the critical bands' concept. An expression for the Bark rate is [56]

$$\begin{aligned} \mathrm{AUD}_{\mathrm{Bark}} &= F_{\mathrm{Bark}}(\xi) \\ &= 13 \arctan(0.00076\xi) + 3.5 \arctan(\xi/7500)^2 \,. \end{aligned} \tag{1}$$

$\mathrm{AUD}_{\mathrm{Bark}}$ corresponds to the critical band rate expressed in Barks. The corresponding bandwidth in Hz is

$$BW_{\mathrm{Bark}} = 25 + 75 \left(1 + 1.4 \times 10^{-6} \xi^2\right)^{0.69} \,. \tag{2}$$

2.2.2. *The ERB Scale.* Follows the same concept as the Bark scale but results from a different set of experiments (see e.g. [33] for a comparison of the two scales and their underlying assumptions). The ERB rate is [18]

$$\mathrm{AUD}_{\mathrm{ERB}} = F_{\mathrm{ERB}}(\xi) = 9.265 \ln\left(1 + \frac{\xi}{228.8455}\right) \tag{3}$$

and the corresponding bandwidth in Hz is

$$BW_{\mathrm{ERB}} = 24.7 + \frac{\xi}{9.265} \quad . \tag{4}$$

Remark: BW_{Bark} and BW_{ERB} are commonly used in psychoacoustics and signal processing to approximate the auditory frequency resolution at low to moderate levels (i.e. 30–70 dB) where the auditory filters' shape remains symmetric and constant. See for example [18, 49] for the variation of BW_{ERB} with level.

2.2.3. *The Mel Scale.* Provides a means to quantify pitch, that is the perceived height of a note, as a function of frequency. A popular formula for the Mel scale is [37]

$$\mathrm{AUD}_{\mathrm{Mel}} = F_{\mathrm{Mel}}(\xi) = 2595 \log_{10}\left(1 + \frac{\xi}{700}\right) \quad . \tag{5}$$

The resulting pitch value has the unit "mel" (as in *mel*ody). Because the Mel scale is not directly related to the auditory filters' concept, it provides no expression for the bandwidth. However, it is common practice to construct Mel FBs with filters linearly distributed on the Mel scale. Their bandwidth is usually set to reach 50% overlap between channels. This is done, for instance, to compute the so-called "Mel-Frequency cepstrum coefficients" (MFCCs) in the fields of speech recognition [44] or music information retrieval [51].

2.3. **Related Work.** A wide variety of tools is available for achieving a perceptually motivated TF transform of a sound. Nonetheless, only a number of these tools allows to (approximately) reconstruct the signal from the transform coefficients. Many analysis-synthesis systems have been proposed that implement gammatone filters in the analysis stage and their time-reversed impulse responses in the synthesis stage (e.g. [20, 22, 27, 47, 54]). This setting implies that the frequency response of the gammatone FB has an all-pass characteristic and features no ripple (equivalently in the frame context, that the system is tight, see Appendix A) while in practice it does not. A reason for that is that gammatone FBs usually consider only a limited range of frequencies (typically in the interval 0.1–4 kHz for speech processing). Therefore, such systems only achieve an approximate reconstruction, still the audio quality of the reconstruction is good provided a rather high density of filters is used [20, 27, 47]. Moreover, commonly-used 4-th order gammatone filters (an order of 4 allows to best approximate the auditory filters' shape at low to moderate levels [39]) do not feature a steep decay in the frequency domain and, therefore, might not have a good aliasing suppression property (this is assessed in Sec. 4).

Other popular tools like auditory models are motivated by the idea to replicate the nonlinear processing in the auditory system (e.g. [24, 28, 36, 55]). Such models are useful to improve our knowledge about the auditory system but they are generally intended for signal analysis only. Thus, they do not allow for signal reconstruction. Note that the approach proposed in [24] does feature a re-synthesis option. This analysis-synthesis system implements compressive gammachirp filters. Nevertheless, because the synthesis stage uses the time-reversed gammachirp impulse responses, the reconstruction is only approximate and a rather large density of filters is required to achieve a good quality, as for gammatone filters.

Other tools include FB constructions aimed at mimicking the auditory frequency resolution for audio signal processing purposes. For instance, wavelet and constant-Q FBs (e.g. [23, 43, 52]) are often used but they mismatch the auditory frequency resolution at low frequencies (< 2 kHz, see e.g. (4)), where they also unnecessarily feature a large number of filters, and do not always achieve perfect reconstruction. Other approaches include FBs made of blocks of uniform frequency resolution FBs [6, 9]. The approach in [6] only roughly approximates the ERB scale and, being targeted at audio coding, is not invertible. The approach in [9] is designed to maximize resistance to sub-channel processing (i.e. filters with a high attenuation outside the passband) and allows for nearly perfect reconstruction for redundancies ≥ 2. However, since one has to properly design the windows and transition filters so that the desired FB properties hold, the global FB design turns out to be complex.

3. Proposed Approach: The AUDlet Filter Bank

The present section introduces a perfect reconstruction oversampled FB that approximates the auditory frequency resolution and provides adaptable redundancy. To allow for a simple and flexible FB design, the FB construction is directly performed in the frequency domain and any compactly supported (e.g. FIR) window is an eligible filter's shape. It has to be considered, however, that the proposed concept is a *linear* FB and, as such, it does not constitute an attempt to reproduce the actual *non*linear auditory filtering. This is discussed below.

3.1. Notation and Definitions.

In the following, we consider signals in $\ell_2(\mathbb{Z})$ sampled at the frequency ξ_s. The inner product of two signals x, y is $\langle x, y\rangle = \sum_n x[n] \cdot y[n]$ and the energy of a signal is defined from the inner product as $||x|| = \langle x, x\rangle$. We denote the z-transform by $\mathcal{Z} : x[n] \mapsto X(z)$. By setting $z = e^{2i\pi\xi}$ for $\xi \in \mathbb{T} := \mathbb{R}/\mathbb{Z}$, the z-transform equals the discrete-time Fourier transform (DTFT). Throughout the paper, bold italic letters indicate matrices (upper case), e.g. $\boldsymbol{G}$, and vectors (lower case), e.g. $\boldsymbol{h}$.

The proposed AUDlet FB has a general non-uniform structure as presented in Fig. 5a with analysis filters $H_k(z)$, synthesis filters $G_k(z)$, and downsampling and upsampling factors d_k. Since we consider signals in $\mathbb{R}$ we deal with symmetric DTFTs, which allows us to process only the positive-frequency range. Therefore, the letter K denotes the number of filters in the frequency range $[\xi_{\min}, \xi_{\max}] \cap [0, \xi_s/2[$, where $\xi_{\min} \geq 0$ to $\xi_{\max} \leq \xi_s/2$ and $\xi_s/2$ is the Nyquist frequency. If $\xi_{\min} > 0$, K includes an additional filter at the zero frequency. Furthermore, another filter is always positioned at the Nyquist frequency. Assuming $\xi_{\min} = 0$ and $\xi_{\max} = \xi_s/2$ for the rest of this manuscript, this implies that all non-uniform FBs treated below feature $K + 1$ filters in total and their redundancy is defined as $R = d_0^{-1} + 2\sum_{k=1}^{K-1} d_k^{-1} + d_K^{-1}$, since coefficients in the 1st to K-th subbands are complex-valued.

We describe below the analysis and synthesis stages of the AUDlet FB and specify the perfect reconstruction conditions for different sets of downsampling factors d_k's. To allow for FB inversion, the analysis FBs described below are constructed such that they always form a frame, i.e. with sufficiently small downsampling factors. For results regarding suitable choices of downsampling factors and references regarding the inversion (*perfect reconstruction conditions*) of non-uniform FBs we refer to the appendix. By default, the algorithms referenced in this manuscript (see Sec. 3.4) automatically determine suitable downsampling factors d_k's.

3.2. **Analysis Filter Bank.** The AUDlet filters H_k's, $k \in \{0, \dots, K\}$ are constructed in the frequency domain by

$$H_k(e^{2i\pi\xi}) = \Gamma_k^{-\frac{1}{2}} w\left(\frac{\xi - \xi_k}{\Gamma_k}\right) \tag{6}$$

where $w(\xi)$ is assumed to be a prototype filter's shape with bandwidth 1 and center frequency 0. This implies that the shape factor Γ_k controls the effective bandwidth of H_k and ξ_k determines its center frequency. The factor $\Gamma_k^{-\frac{1}{2}}$ ensures that all filters (i.e. $\forall\ \xi_k$) have the same energy. To obtain filters equidistantly spaced on a perceptual frequency scale, the sets $\{\xi_k\}$ and $\{\Gamma_k\}$ are calculated using the corresponding $\mathrm{AUD}_{\mathrm{scale}}$ and BW_{scale} formulas. For instance, linearly distributing K filters from $\mathrm{AUD}_{\mathrm{ERB}_{\min}} = F_{\mathrm{ERB}}(\xi_{\min})$ to $\mathrm{AUD}_{\mathrm{ERB}_{\max}} = F_{\mathrm{ERB}}(\xi_{\max})$ with a density of V filters per ERB leads to an ERB step $\mathrm{AUD}_{\mathrm{ERB}_k} = \mathrm{AUD}_{\mathrm{ERB}_{\min}} + k/V$. Then $\xi_k = F_{\mathrm{ERB}}^{-1}(\mathrm{AUD}_{\mathrm{ERB}_k})$ and $\Gamma_k = BW_{\mathrm{ERB}}(\xi_k)$. Overall, the resolution of the analysis is given by two parameters: $K = V\left(\mathrm{AUD}_{\mathrm{ERB}_{\max}} - \mathrm{AUD}_{\mathrm{ERB}_{\min}}\right)$ and the set of downsampling factors $\{d_k\}$. An analogous process yields an FB adapted to any frequency scale.

3.3. **Synthesis.** In general, the existence of a non-uniform dual FB having the same number of filters G_k's and upsampling factors d_k's as the non-uniform analysis FB cannot be guaranteed. Therefore, we use three different approaches to compute the action of the AUDlet synthesis FB:

(i) For band-limited filters with sufficiently dense sampling, dual synthesis filters can be explicitly and efficiently computed. Synthesis is then accomplished by a standard non-uniform FB synthesis algorithm. The formal conditions for this setting are given in Thm 1 in the appendix. The dual FB is computed by (19) also given there.

(ii) If the conditions of Thm 1 are violated but $\sum_{k=0}^{K} q_k$ is small enough, then the equivalent uniform FB for H_k, d_k, $k \in \{0, \dots, K\}$, is constructed as described in the appendix, see Fig. 5b. A dual FB can be easily obtained using standard algorithms for the computation of dual uniform FBs.

(iii) If the number of channels in the equivalent uniform FB is too large, the computation and storage of the dual FB become unfeasible. In such cases, the action of the canonical dual FB is computed using a conjugate gradients (CG) algorithm. Iterative synthesis via CG benefits from the fact that although the number of iterations necessary to achieve the desired precision depends on the actual frame bound ratio of the analysis FB, it does not require explicit estimates of the frame bounds as opposed to other iterative approaches like the classical frame algorithm [19]. Furthermore, since each iteration computes the analysis followed by synthesis with the filters H_k's, see (21), the algorithm's complexity is independent of the structure of the dual FB. Additionally, we showed in [35] that using a preconditioner often drastically reduces the number of iterations required to achieve a certain precision.

3.4. **Implementation.** For the implementation we consider finite-length sequences in $\mathbb{C}^L$, $L \in \mathbb{N}$. For the extension of the results in Appendix A to finite-length sequences we refer to [14]. We provide code for performing an AUDlet analysis/synthesis as part of the Matlab/Octave "LTFAT" toolbox [40] available at

`http://ltfat.sourceforge.net/`. The analysis filters are generated by the function `audfilters`. The function allows to construct at will uniform or non-uniform AUDlet FBs with integer or rational downsampling factors,[2] thus offering flexibility in FB design. The desired number of channels can be set by specifying either K or V. Using the block-processing framework proposed in [23], a real-time AUDlet analysis is also possible in LTFAT with block-stream processing [41]. As for the prototype window w, the function `audfilters` uses by default a Hann window, but any FIR window can be chosen. In Sec. 4 we present results obtained with different window types. The synthesis FB is computed by the function `filterbankdual` for the cases (i) and (ii) mentioned above. For case (iii), the pseudo code is presented on the Web page associated with [35]. Analysis and synthesis are finally performed by the functions `filterbank` and `ifilterbank`, respectively.

4. Experiments

To illustrate the properties and signal processing capabilities of the AUDlet FB, we present in this section the results from three experiments. The first experiment is a direct comparison between the proposed framework and a classic linear auditory FB, namely the gammatone FB, in terms of FB response, reconstruction error, aliasing suppression, and signal representation. The effect of the prototype filter's shape w on the aliasing suppression property of the FB is also investigated. The two follow-up experiments are exemplar applications of the AUDlet FB to audio signal processing, namely a source separation and a speech de-noising experiment. For demonstration purposes all FBs were adapted to the ERB scale. The resulting AUDlet FB is thus called "ERBlet FB" in the following. Results on the performance of an ERBlet iterative reconstruction using CG can be found in [35].

4.1. Filter Bank Settings.

All experiments presented below feature non-uniform ERBlet and gammatone FBs. The discrete-time impulse responses of the gammatone filters were calculated by sampling and windowing the complex continuous-time gammatone IIR

$$h_{\mathrm{gt,k}}(t) = \alpha_k t^{\gamma-1} e^{2\pi t(i\xi_k - \lambda_k)} \quad t \geq 0, k \in \{0, \ldots, K\} \tag{7}$$

where ξ is the filter center frequency, γ is the gammatone filter order, $\lambda_k = \beta\,\mathrm{ERB}(\xi_k)$ determines the filter bandwidth and α_k is a normalization factor that constraints all filters to have the same energy. We chose $\gamma = 4$ and $\beta = 1.019$ to obtain a gammatone FB adapted to human auditory perception [39]. This FIR filter design allows for straightforward implementation but it requires rather long impulse responses to correctly approximate the filter responses at low frequencies. We used a length of 6000 samples. The gammatone synthesis FB consists of the synthesis filters $g_{\mathrm{gt,k}}[n] = \overline{h_{\mathrm{gt,k}}}[-n]$, where the bar denotes the complex conjugate, and upsampling factors d_k's as the analysis FB. Using time-reversed versions of the analysis filters for synthesis might not be the best setting in terms of reconstruction error (this is discussed below) but this is the most common use of gammatone FBs

[2] Although the results stated in Appendix A are valid only for d_k's $\in \mathbb{Z}$, rational downsampling factors can be achieved in the time domain by properly combining upsamplers and downsamplers (e.g. [25]). In LTFAT the sampling rate changes are directly performed in the frequency domain by periodizing and folding the $Y_k(z)$'s, then performing an inverse DFT [43]. This technique allows to achieve rational downsampling factors at low computational costs.

in audio applications (e.g. [20,22,27,47,54]). We therefore stick to this setting. Additionally, note that in contrast to previous band-limited gammatone FB designs, our gammatone FB covers the full range of frequencies (i.e. from 0 to the Nyquist frequency).

To achieve a fair comparison between the ERBlet and gammatone systems, both FBs were configured identically, specifically with the same spectral and temporal resolutions. In Experiment 1, a spectral resolution of $V = 1$ filter/ERB was chosen to show results with a small density of filters. All filters were 1-ERB wide. In Experiments 2 and 3, a spectral resolution of $V = 6$ filters/ERB was chosen to achieve good signal processing performance. Since applications often require a finer resolution than that provided by 1-ERB-wide filters, especially at high frequencies where the ERBs are large, we set the bandwidths of the ERBlet and gammatone filters to one sixth of an ERB (i.e. $\Gamma_k = \mathrm{ERB}(\xi_k)/6$ and $\beta = 1.019/6$) in those experiments. In this setting the FBs are only partly perceptually motivated, though. All ERBlet calculations were performed using a Hann window as w, unless otherwise stated. A set of integer downsampling factors was generated that satisfies Thm 1 for the Hann ERBlet. Let R_i denote the redundancy of the resulting painless system. We then evaluated the FB performances for four redundancies $R = redfac \times R_i$ with $redfac = 0.38$, $1/2$, 1.0 and 2.0. For $redfac = 0.38$ and $1/2$, Thm 1 is violated and the ERBlet synthesis is done using CG.

4.2. **ERBlet vs. Gammatone FB.** Fig. 1 shows the magnitudes of the ERBlet (solid line) and gammatone (dashed line) FB responses in the frequency range from zero to the Nyquist frequency. While the ERBlet FB response is rather flat across all the passband, the gammatone FB response features significant ripples. These ripples are likely to affect the reconstruction property of the gammatone FB. As is easily seen, defining synthesis filters that are time-reversed versions of the analysis filters infers that the FB response is constant over the full frequency range, cp. (19) in the appendix, noting that $\mathcal{H}_0(\xi)$ corresponds to the FB response. Therefore, the results in Fig. 1 indicate that the gammatone FB in this particular setting is not a perfect reconstruction system. Accordingly, the relative reconstruction errors for the gammatone and ERBlet FBs for the four redundancy factors are listed in Table 1. While the ERBlet scheme using the proposed methods always achieves perfect reconstruction up to numerical precision, using gammatone filters in the analysis and (time-reversed) in the synthesis step generates a relative error of about 10^{-1}. A similar reconstruction error was reported in [47] using FIR gammatone filters and about 1 filter per ERB. Noteworthy, the present gammatone reconstructions for the two smallest redundancy factors featured audible distortions that are likely due to the ripples in the gammatone FB response and the gammatone filters' weak aliasing suppression.

To assess the aliasing suppression capability of each FB, the magnitude responses (in dB) of the gammatone (gray dashed) and Hann ERBlet filters (black solid) for channel $k = 28$ are plotted in Fig. 2 in the frequency range $[0; \xi_s/2]$. The magnitude responses of a Gaussian (black dashed) and a roex ERBlet filter (gray solid) are also shown. For the roex variant, the prototype filter w was a symmetric roex(p,r) defined by its frequency response [49]

$$H_{\mathrm{roex},k}(e^{2i\pi\xi}) = (1-r)(1+p_k|\xi-\xi_k|/\xi_k)e^{-p_k|\xi-\xi_k|/\xi_k} + r, \tag{8}$$

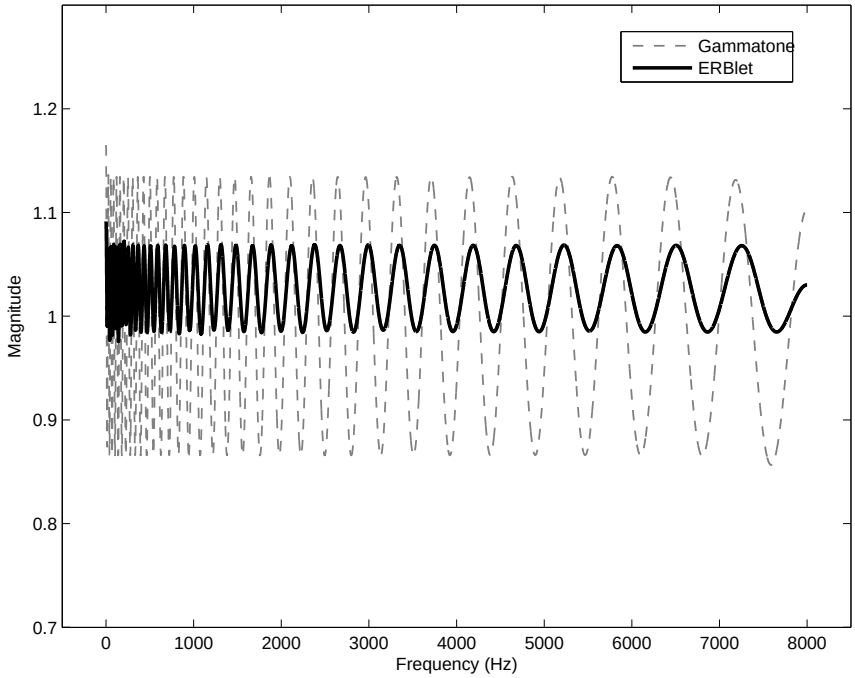

FIGURE 1. Magnitudes of the FB responses in the frequency range $[0; \xi_s/2]$ for the gammatone (dashed line) and ERBlet (solid line) FBs using $V = 1$ ($K = 34$).

$k \in \{0, \ldots, K\}$ with $r = 0$.[3] The parameter p_k was tuned to obtain a bandwidth of 1 ERB. The roex filter is an analogue of the gammatone filter but defined in the frequency domain. It is thus more easily applicable to the AUDlet FB's concept than the gammatone.

To avoid aliasing in channel k the signal must be band-limited to ξ_s/d_k around ξ_k. In the example illustrated in Fig. 2 $d_{28} = 14$, $\xi_{28} \approx 3.7$ kHz and $\xi_s/d_{28} \approx 1200$ Hz, that is the filters must have a high attenuation for frequencies outside the range $[3.1; 4.3\,\text{kHz}]$. While all filters have a similar attenuation for frequencies between 3.4 and 4 kHz, the attenuations of the Hann and Gaussian ERBlet filters are superior to those of the roex and gammatone filters for other frequencies. This indicates that the ERBlet FB has a better resistance to sub-channel processing than the gammatone FB. This is verified in the two follow-up experiments.

Finally, the ERBlet (a) and gammatone (b) analyses of speech signals are represented in Fig. 3 for $redfac = 2$. This experiment was performed on a female speech signal sampled at 16 kHz taken from the TIMIT database [17]. To better represent the harmonics, we chose $V = 6$. It can be seen that the two signal representations are very similar over the whole TF plane.

[3]This setting can be easily adjusted to a parallel roex filter that better represents the auditory filters' shape than a single roex(p,r) [49]. However, simulations showed that using a parallel roex only deteriorates the roex' aliasing suppression outside the passband. Consequently, there is no significant impact on the reconstruction error at moderate redundancies but reconstruction error increases drastically at low redundancies, similar to the gammatone FB.

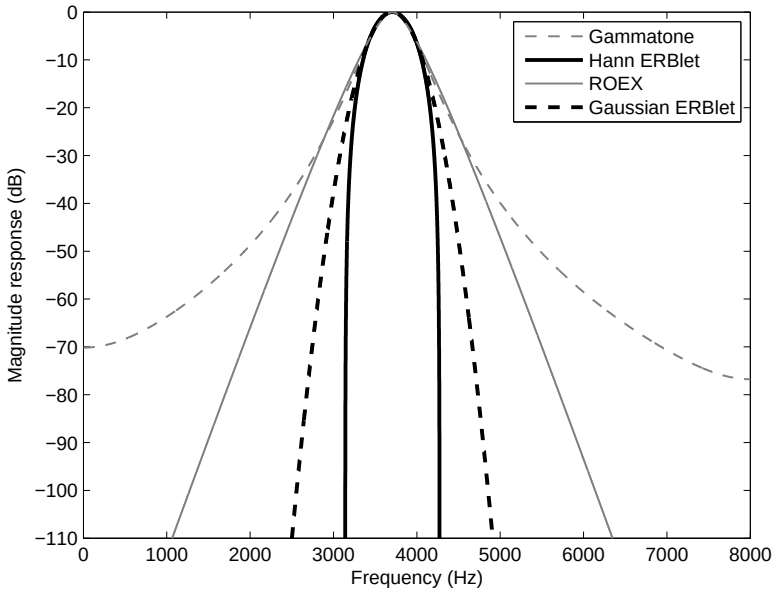

FIGURE 2. Magnitude responses (in dB) of $H_{\mathrm{gt},28}(\xi)$ (dashed gray) and $H_{28}(\xi)$ (solid black) in the frequency range $[0; \xi_s/2]$. Additionally, a roex (solid gray) and a Gaussian filter (dashed black) with the same center frequency and ERB are shown.

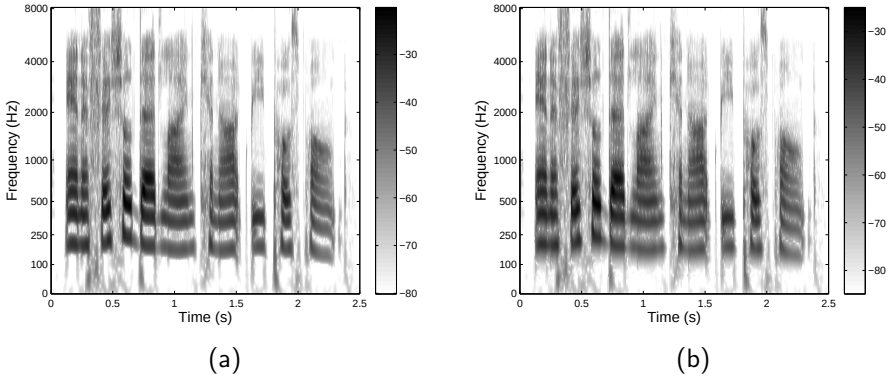

FIGURE 3. Analyses of a speech signal by (a) the ERBlet FB and (b) the gammatone FB using $V = 6$ ($K = 201$) and $redfac = 2$.

4.3. **Separation by Masking.** In this experiment, we attempt the separation of a musical source from a vocal and music mixture through a simple masking procedure using a binary mask [54]. We evaluate the separation using classical signal-to-noise ratio (SNR, in dB) but also the following measures proposed in [53]

TABLE 1. Relative reconstruction errors for the gammatone, ROEX and Hann ERBlet FBs using $V = 1$ for various redundancy factors. The corresponding actual redundancies R's are also indicated.

$redfac$	0.38	1/2	1	2
R	1.13	1.48	3.04	6.18
gammatone FB	0.55	0.34	0.10	0.10
ROEX FB	0.67	0.43	0.12	0.11
ERBlet FB	$1\text{x}10^{-14}$	$4\text{x}10^{-15}$	$5\text{x}10^{-16}$	$5\text{x}10^{-16}$

TABLE 2. Quality measures for the separation experiment with low redundancy ($redfac = .38, R = 1.06$)

	SDR	SIR	SAR	SNR
ERBlet separated voice	11.51	16.35	13.33	11.43
ERBlet separated music	7.16	16.85	7.75	6.53
Gammatone sep. voice	3.87	15.57	4.29	3.73
Gammatone sep. music	−1.65	12.44	−1.24	−1.17

for determining separation performance: signal-to-distortion ratio (SDR), signal-to-interference ratio (SIR) and signal-to-artifact ratio (SAR), all in dB. For their exact definition, please refer to [53].

The binary mask was created by a combination of thresholding the ERBlet spectrogram of the vocal source and manual editing in an image processing software. Separation is performed by (i) analyzing the mixture with an ERBlet (gammatone) FB, (ii) applying the mask to the coefficients by point-wise multiplication and (iii) synthesizing with the dual FB (the time-reversed analysis FB).[4] Fig. 4 shows the mixture (a), ground truth (b), separated signals (c,d), and the separation mask (e) in a high redundancy setup ($redfac = 2$). For subplots (c) and (d), the separated signal was re-analyzed using the analysis FB used in the processing step.

The separation results for low, medium and high redundancy setups are listed in Tables 2–4, respectively. It can be seen that the ERBlet FB slightly outperforms the gammatone FB in practically every measure, with the possible exception of SIR on the vocal source, where the results are still roughly equivalent. Note that the vocal source is not the target signal and the mask was designed to remove the voice from the music, i.e. interference in the separated voice signal is preferred over interference in the separated music source. One can also see that the performance differences between gammatone and ERBlet FBs increase with decreasing redundancy, further illustrating the superior processing stability and reconstruction quality of the ERBlet FB, in particular at low-redundancy setups.

[4]Such an approach is often used, for example, in computational auditory scene analysis [54] and is known in mathematical signal processing as a frame multiplier [46].

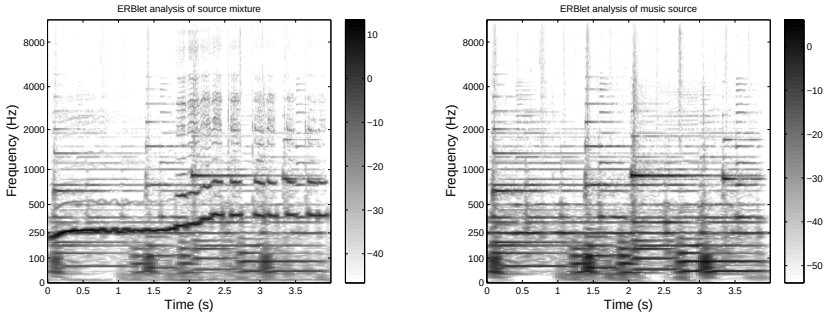

(a) Signal mixture

(b) Ground truth signal

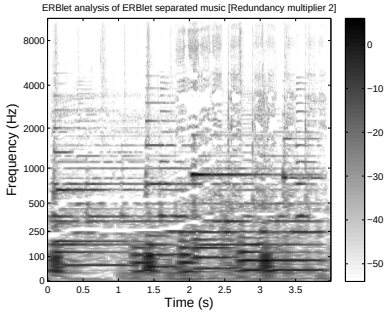

(c) ERBlet separated signal

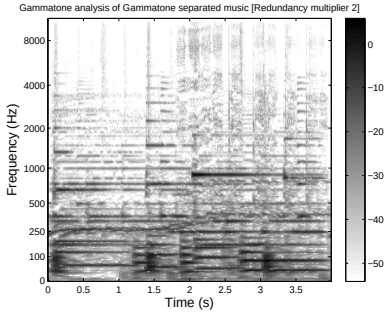

(d) Gammatone separated signal

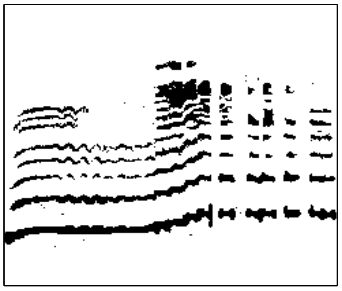

(e) Binary separation mask

FIGURE 4. Inputs (a-b) and separation results for (c) the ERBlet and (d) the gammatone FB. Note how the gammatone separation shows considerably more residual energy from the vocal source than the ERBlet separation, particularly in the first section of the signal. This is a consequence of the gammatone filters' inferior TF concentration. (e) Binary mask used in the separation process, where black pixels represent the masks 0 entries.

TABLE 3. Quality measures for the separation experiment with medium redundancy ($redfac = 1, R = 2.78$)

	SDR	SIR	SAR	SNR
ERBlet separated voice	13.85	16.56	17.28	13.65
ERBlet separated music	10.92	21.16	11.39	8.75
Gammatone sep. voice	12.73	16.61	15.11	12.55
Gammatone sep. music	8.88	20.44	9.24	7.65

TABLE 4. Quality measures for the separation experiment with high redundancy ($redfac = 2, R = 5.59$)

	SDR	SIR	SAR	SNR
ERBlet separated voice	13.89	16.66	17.25	13.69
ERBlet separated music	10.89	21.10	11.36	8.79
Gammatone sep. voice	12.88	16.65	15.33	12.67
Gammatone sep. music	9.07	20.41	9.44	7.77

4.4. **Speech De-Noising.** In this third and last experiment, we perform a de-noising task in the transform domain, compare to [3]. Specifically, we apply channel-wise the soft-thresholding function [12]

$$y_{k_{\text{thr}}} = \text{sgn}(y_k)(|y_k| - \eta)_+$$

to the sub-band components $y_k[n]$ where η is the threshold value. The de-noised signal is then obtained by applying the synthesis FB to the modified components $y_{k_{\text{thr}}}$. The audio material consisted of a male and a female extract taken from the TIMIT database ($\xi_s = 16$ kHz) [17]. The signals were corrupted by Gaussian white noises of different powers. Denote by σ the standard deviation (power) of the corrupting noise. We set the threshold parameter $\eta = \sigma$. The de-noising performance is evaluated by two measures: the SNR and segmental SNR (segSNR). The segSNR measures were computed as in [21] using 32-ms frames and limited to the range of -10 to 35 dB. Table 5 compares the SNR (top) and segSNR (bottom) of the ERBlet and gammatone FBs for various noise powers and redundancy factors. It can be seen that the ERBlet systematically achieves higher SNR and segSNR than the gammatone. At low redundancy, the ERBlet improves the SNR by up to 9.5 dB (average gain = 5 dB) and the segSNR by up to 5 dB (average gain = 2.7 dB). However, the differences in SNR and segSNR are very small ($\leq$1 dB) at medium and high redundancies. Noteworthy, running this experiment without sub-sampling (i.e. $d_k = 1\ \forall k$) resulted in the same SNR and segSNR values as with $redfac = 2$ for both FBs, that is it did not improve performance. Overall, the better performance of the ERBlet again illustrates its perfect reconstruction and good resistance to sub-channel processing as compared to the gammatone FB, especially at low redundancy.

TABLE 5. Comparison of the SNR (top) and segSNR values (bottom) of the ERBlet and gammatone FBs for two types of speech signals with different noise powers. For each signal, the corrupting noise power is indicated as the input SNR in dB.

Signal & input SNR	$redfac = .38$ $R = 1.1$		$redfac = 1$ $R = 2.9$		$redfac = 2$ $R = 5.9$	
	GFB	ERBlet	GFB	ERBlet	GFB	ERBlet
Male -5 dB	1.64	3.33	4.21	4.29	4.24	4.29
Male 0 dB	3.73	7.28	7.91	8.10	7.94	8.11
Male 10 dB	5.13	14.67	14.12	15.35	14.22	15.35
Female -5 dB	1.47	3.26	4.27	4.33	4.31	4.33
Female 0 dB	3.49	7.18	7.95	8.14	8.00	8.14
Female 10 dB	4.51	14.31	13.95	15.11	14.08	15.12

Signal	$redfac = .38$		$redfac = 1$		$redfac = 2$	
& input SNR	GFB	ERBlet	GFB	ERBlet	GFB	ERBlet
Male -5 dB	-3.55	-2.50	-1.92	-1.85	-1.90	-1.85
Male 0 dB	-1.84	0.22	0.69	0.86	0.71	0.86
Male 10 dB	0.14	5.71	5.23	6.26	5.29	6.26
Female -5 dB	-4.18	-3.25	-2.65	-2.61	-2.63	-2.61
Female 0 dB	-2.58	-0.79	-0.27	-0.13	-0.24	-0.13
Female 10 dB	-1.02	3.93	3.79	4.55	3.88	4.56

5. SUMMARY AND CONCLUDING REMARKS

The construction of an oversampled perfect reconstruction FB with filters distributed on a perceptual frequency scale has been presented. The resulting perceptually motivated FB is named "AUDlet FB". The FB design is directly performed in the frequency domain and allows for various filter shapes, uniform or non-uniform setting, and large downsampling factors. For redundancies ≥ 3 (i.e. ensuring a painless system), the synthesis (dual) filters are explicitly computed. For lower redundancies, an iterative algorithm is used to compute the action of the dual FB. The TF resolution and redundancy of the FB are adaptable without affecting its perfect reconstruction property down to redundancies close to 1. Overall, the proposed system provides a simple and efficient FB design that is highly suitable for audio applications that require an analysis-synthesis framework. We provide an implementation of the AUDlet FB in the free Matlab/Octave toolbox LTFAT.

An experiment compared the AUDlet to a linear auditory FB that is widely used in audio applications, namely the gammatone FB. The results showed the better performance of the AUDlet FB with respect to the gammatone FB in terms of reconstruction error and resistance to sub-channel signal processing, especially at low redundancies. Two additional experiments demonstrated the utility of the AUDlet FB as an audio processing tool.

The proposed concept is a linear FB and, as such, does not constitute a realistic auditory filter's model, as proposed for instance in [24, 28]. In particular, we do

not consider nonlinearities due to varying sound pressure levels (SPLs). Both of the cited approaches (respectively the dual-resonance nonlinear [28] and dynamic compressive gammachirp FBs [24]) feature a *linear* filter in the first stage and the nonlinearities are added subsequently. It is thus conceivable that a similar nonlinear FB construction be achieved using an AUDlet FB, for instance by adding a compressive nonlinearity subsequent to the AUDlet filters. Nevertheless, this is likely to alter the stability and perfect reconstruction property of the analysis-synthesis system, especially if sub-channel processing is performed. Considering that in many applications the SPL is unknown in the signal chain (the SPL actually depends on the final listening volume), using level- *in*dependent filters is the most conservative course of action and may suffice in most cases.

To further reduce the redundancy of the AUDlet representation and improve its perceptual relevance, future work includes introducing perceptual sparsity in the transform domain. Specifically, based on the perceptual irrelevance filter proposed in [5] and recent data on auditory TF masking [34], a binary mask will be computed and applied to the sub-channel coefficients in order to re-synthesize only the audible TF components. Furthermore, future work will focus on how to combine the AUDlet FB and knowledge of TF masking to possibly improve audio codecs. For a first approach on how to adapt the ERBlet FB for audio coding see [11].

Acknowledgment

The authors would like to thank Damián Marelli for insightful discussions and help on the theoretical development on non-uniform FBs.

Appendix A. Invertibility of non-uniform FBs

Consider a non-uniform FB structure with downsampling and upsampling factors $d_k \in \mathbb{Z}$ as depicted in Fig. 5a. In this appendix, we denote by $W_N = e^{2i\pi/N}$ the Nth root of unity.

A.1. **Perfect Reconstruction Conditions.** The sub-band components of the system represented in Fig. 5a are given in the discrete-time domain by

$$y_k[n] = \downarrow_{d_k} \{h_k * x\}[n]. \tag{9}$$

The output signal is $\tilde{x}[n] = \sum_{k=0}^{K} (g_k * \uparrow_{d_k} \{y_k\})[n]$. Such an analysis-synthesis system provides perfect reconstruction if $\tilde{x}[n] = x[n]$ (up to a delay factor). When $d_k = D \,\forall\, k \in \{0 \dots K\}$ the system in Fig. 5a results in a uniform FB. The perfect reconstruction conditions for *uniform* FBs have been largely treated in the literature (see e.g. [25, 50]). To treat the *non-uniform* case, one possibility is to decompose the non-uniform system into a larger equivalent uniform system [1, 25], as shown in Fig. 5b. Denote $D = \mathrm{lcm}(d_k : k \in \{0, \cdots, K\})$ and $q_k = D/d_k$. Each k-th channel of the non-uniform system is decomposed into q_k channels in the equivalent uniform system, which then features $\sum_{k=0}^{k=K} q_k$ channels in total with the downsampling factor D in all channels. footnoteNote that for a maximally decimated non-uniform FB, i.e. when $\sum_k 1/d_k = 1$, the equivalent uniform FB features D channels. Note that the filters in Fig. 5b are various delayed versions of those in Fig. 5a. The sub-band components for $l \in \{0, \cdots, q_k - 1\}$ in Fig. 5b can be expressed in the

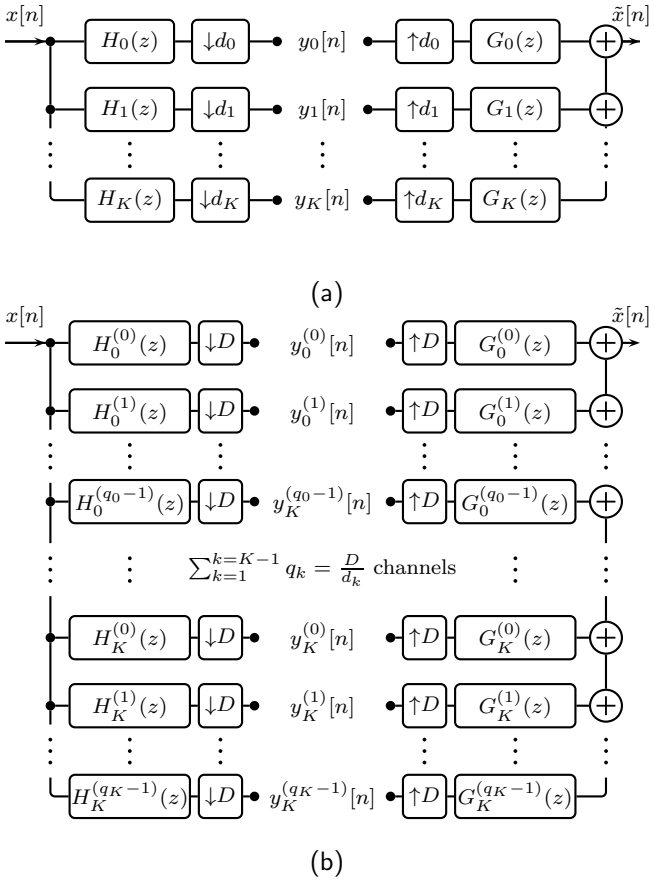

FIGURE 5. (a) General structure of a non-uniform analysis-synthesis FB and (b) an equivalent uniform FB [1]. The terms $H_k^{(l)}$ and $G_k^{(l)}$ in (b) correspond to the z-transforms of the terms $h_k^{(l)}$ and $g_k^{(l)}$ defined in (10) and (11), respectively.

discrete-time domain as

$$\begin{aligned} y_k^{(l)}[n] &= y_k[nq_k - l] \\ &= \downarrow_D \{\underbrace{h_k * \delta_{ld_k}}_{:=h_k^{(l)}} * x\}[n]. \end{aligned} \tag{10}$$

Grouping the q_k sub-band components resulting from the k-th sub-band yields $y_k[n] = \sum_{l=0}^{q_k-1} \uparrow_{q_k} \left\{y_k^{(l)}\right\}[n+l]$ and the output signal of the equivalent uniform FB can be written as

$$\tilde{x}[n] = \sum_{k=0}^{K} \sum_{l=0}^{q_k-1} \left(\underbrace{g_k * \delta_{-ld_k}}_{:=g_k^{(l)}} * \uparrow_D \left\{y_k^{(l)}\right\} \right)[n]. \tag{11}$$

These results can be written in the frequency domain by simply taking the z-transform. The sub-band components *after* upsampling by D yield

$$\begin{aligned} Y_k^{(l)}\left(z^D\right) &= \frac{1}{D}\sum_{j=0}^{D-1} H_k^{(l)}\left(W_D^j z\right) X\left(W_D^j z\right) \\ &= \frac{1}{D}\sum_{j=0}^{D-1} W_D^{-jld_k} z^{-ld_k} \underbrace{H_k\left(W_D^j z\right) X\left(W_D^j z\right)}_{\text{for } j\neq 0\text{: aliasing terms}} \end{aligned}$$

and the output finally yields

$$\begin{aligned} \widetilde{X}(z) &= \sum_{k=0}^{K}\sum_{l=0}^{q_k-1} G_k^{(l)}(z) Y_k^{(l)}\left(z^D\right) \\ &= \frac{1}{D}\sum_{k=0}^{K}\sum_{j=0}^{D-1} A_{k,j} G_k(z) H_k\left(W_D^j z\right) X\left(W_D^j z\right) \end{aligned} \tag{12}$$

where

$$A_{k,j} := \sum_{l=0}^{q_k-1} W_D^{-jld_k} = \sum_{l=0}^{q_k-1} e^{-2i\pi jl/q_k} = \begin{cases} q_k & \text{if } j \text{ is a} \\ & \text{multiple of } q_k \\ 0 & \text{otherwise.} \end{cases}$$

Equation (12) can be formulated as a matrix multiplication

$$\widetilde{X}(z) = \frac{1}{D}\left[X(W_D^0 z)\cdots X(W_D^{D-1} z)\right] \boldsymbol{H}(z)\boldsymbol{G}(z) \tag{13}$$

where $\boldsymbol{G}(z) := [G_0(z), \ldots, G_K(z)]^T$ and the $D\times K$ alias cancellation matrix $\boldsymbol{H}(z) = [\boldsymbol{h}_0(z)\cdots\boldsymbol{h}_K(z)]$, cf. [1], with:

$$\boldsymbol{h}_k(z) = q_k \begin{bmatrix} \boldsymbol{h}_k'(z) \\ \boldsymbol{h}_k'\left(W_D^{q_k} z\right) \\ \vdots \\ \boldsymbol{h}_k'\left(W_D^{(d_k-1)q_k} z\right) \end{bmatrix}$$

and

$$\boldsymbol{h}_k'(z) = \Big[H_k(z)\ \underbrace{0\cdots 0}_{q_k-1 \text{ zeros}}\ \Big]^T.$$

The perfect reconstruction condition then reduces to

$$\boldsymbol{H}(z)\,\boldsymbol{G}(z) = [D\ 0\cdots 0\ ]^T, \tag{14}$$

which means that all aliasing terms have to be canceled by the synthesis filters. Equation (14) is useful to determine whether a complete FB provides perfect reconstruction. It may however fail to provide straightforward or efficient ways to find, given fixed analysis parameters h_k's and d_k's, fitting synthesis filters and downsampling factors, although it can sometimes be used to determine whether such a system even exists.

A.2. **Connection to Frame Theory.** Since our construction of perceptually motivated FBs emphasizes stable reconstruction from the FB coefficients, it seems worthwhile to mention that inversion of non-uniform FBs can also be investigated using frame theory, the mathematical theory of stable, redundant spanning sets of functions. For uniform FBs, this connection has been explored in depth (e.g. [3, 8, 10, 14]). For non-uniform FBs, the connections have, to our knowledge, not been stated explicitly in the literature although implicitly used in recent work (e.g. [23], [7]). A frame over the space of finite energy sequences $\ell_2(\mathbb{Z})$ is a (possibly redundant) family of functions spanning the space in a stable fashion, in the sense of inequality (15) below. The central observation linking FBs to frames is that

$$y_k[n] = {\downarrow_{d_k}} \{h_k * x\}[n] = \langle x, \overline{h_k}[nd_k - \cdot]\rangle.$$

Hence, the FB coefficients with respect to the filters h_k and downsampling factors d_k equal the frame coefficients of the system $\left(\overline{h_k}[nd_k - \cdot]\right)_{k,n}$. As a consequence, the FB allows for numerically stable perfect reconstruction if and only if $0 < A \leq B < \infty$ exist such that

$$A\|x\|^2 \leq \sum_k \|y_k\|^2 \leq B\|x\|^2, \text{ for all } x \in \ell^2(\mathbb{Z}) \tag{15}$$

where A and B are respectively the lower and upper frame bounds of the system.[5]

Therefore, instead of verifying the perfect reconstruction conditions directly, we can equivalently employ techniques from frame theory to determine the inversion of the FB analysis operation and/or an appropriate synthesis system. In particular, a dual (synthesis) frame can be found by applying the inverse of the frame operator $\mathbf{S}$ defined by

$$\mathbf{S}\, x[n] = \sum_{n,k} \langle x, \overline{h_k}[nd_k - \cdot]\rangle h_k[nd_k - \cdot],$$

to each frame element. More precisely, we will use the inherent structure of the frequency domain variant, the matrix Fourier transform [4] of $\mathbf{S}$, $\widehat{\mathbf{S}} = \mathrm{DTFT}\,\mathbf{S}\,\mathrm{DTFT}^{-1}$ of the frame operator. By applying Eq. 13, we can easily see, using $\mathcal{Z}(\overline{h[-\cdot]})(z) = \overline{\mathcal{Z}(h)(1/\overline{z})}$, that the action of $\mathbf{S}_{\widehat{\Phi}}$ is given by

$$\begin{aligned} \widehat{\mathbf{S}}X(z) &= \frac{1}{D}\left[X(W_D^0 z)\cdots X(W_D^{D-1} z)\right] \boldsymbol{H}(z) \begin{bmatrix} \overline{H_0(1/\overline{z})} \\ \vdots \\ \overline{H_K(1/\overline{z})} \end{bmatrix} \\ &= \left[Y_0(z^{d_0})\cdots Y_K(z^{d_K})\right] \left[\overline{H_0(1/\overline{z})}\cdots\overline{H_K(1/\overline{z})}\right]^T. \end{aligned} \tag{16}$$

Defining

$$\begin{bmatrix} \mathcal{H}_0(\xi) \\ \vdots \\ \mathcal{H}_{D-1}(\xi) \end{bmatrix} := \frac{1}{D}\boldsymbol{H}(e^{2i\pi\xi}) \begin{bmatrix} \overline{H_0(e^{2\pi i\xi})} \\ \vdots \\ \overline{H_K(e^{2\pi i\xi})} \end{bmatrix} \tag{17}$$

for $\xi \in \mathbb{T} = \mathbb{R}/\mathbb{Z}$, we obtain the following results, see [2] and [13] for the mathematical context.

[5] Boundedness of h_k for all k is sufficient for the existence of B since the number of channels is finite.

Theorem 1. *If, for every $0 \leq k \leq K$, the filter h_k is band-limited on an interval of length $1/d_k$, i.e. there is $I_k \subseteq \mathbb{T}$ with $|I_k| \leq 1/d_k$ such that $H_k(e^{2i\pi\xi}) = 0$ for almost every $\xi \in \mathbb{T} \setminus I_k$. Then $\mathcal{H}_k$ equals the zero function for $k = 1, \ldots, K$ and the FB comprised of the filters h_k's and downsampling factors d_k's forms a frame if and only if there are A, B such that*

$$0 < A \leq \mathcal{H}_0(\xi) \leq B < \infty, \text{ for a.e. } \xi \in \mathbb{T}. \tag{18}$$

Moreover, a dual FB frame with upsampling factors d_k's is given by the filters g_k's defined by

$$G_k(e^{2i\pi\xi}) = \frac{H_k(e^{2i\pi\xi})}{\mathcal{H}_0(\xi)} \text{ a.e.} \tag{19}$$

Although the implications of Theorem 1 are relatively well understood, frame theory provides a simple generalization that is very useful when combined with more sophisticated inversion techniques.

Theorem 2. *Let h_k's and d_k's for $0 \leq k \leq K$ define an analysis FB. If there are $0 < A_0 \leq B_0 < \infty$ with*

$$A_0 \leq \mathcal{H}_0(\xi) \pm \sum_{n=1}^{D-1} |\mathcal{H}_n(\xi)| \leq B_0, \text{ for a.e. } \xi \in \mathbb{T}, \tag{20}$$

then the FB defined by h_k's and d_k's forms a frame.

Although reconstruction can be implemented by rewriting the FB as a uniform FB and computing the dual uniform FB, this is only feasible if $\sum_k q_k$ is not too large. However, any FB frame admits reconstruction by a *conjugate gradients* (CG) algorithm [19, 48] solving

$$\widehat{\mathbf{S}} X(z) = \sum_{k=0}^{K} Y_k(z^{d_k}) H_k(z). \tag{21}$$

The number of CG steps necessary for convergence depends solely on the condition number of $\widehat{\mathbf{S}}$. Additionally, if the conditions of Thm 2 are satisfied and the second term in (20) is small, then $\widehat{\mathbf{S}}$ is diagonal dominant and its diagonal equals $\mathcal{H}_0$. In this setting, using $\mathcal{H}_0^{-1}$ as a diagonal preconditioner has been shown to further increase convergence speed [4, 19, 35]. This method often allows for efficient inversion even if direct computation of a dual uniform FB is not feasible.

References

1. Sony Akkarakaran and PP Vaidyanathan, *Nonuniform filter banks: new results and open problems*, Beyond wavelets, Studies in Computational Mathematics, vol. 10, Elsevier, 2003, pp. 259–301.
2. P. Balazs, M. Dörfler, N. Holighaus, F. Jaillet, and G. Velasco, *Theory, implementation and applications of nonstationary Gabor frames*, J. Comput. Appl. Math. **236** (2011), no. 6, 1481–1496.
3. P. Balazs, M. Dörfler, M. Kowalski, and B. Torrésani, *Adapted and adaptive linear time-frequency representations: a synthesis point of view*, IEEE Signal Process. Mag. **30** (2013), no. 6, 20–31.
4. P. Balazs, H. G. Feichtinger, M. Hampejs, and G. Kracher, *Double preconditioning for Gabor frames*, IEEE Trans. Signal Process. **54** (2006), no. 12, 4597–4610.
5. P. Balazs, B. Laback, G. Eckel, and W. A. Deutsch, *Time-frequency sparsity by removing perceptually irrelevant components using a simple model of simultaneous masking*, IEEE Trans. Audio, Speech, Language Process. **18** (2010), no. 1, 34–49.

6. F. Baumgarte, *Improved audio coding using a psychoacoustic model based on a cochlear filter bank*, IEEE Speech Audio Process. **10** (2002), no. 7, 495–503.
7. I Bayram, *An analytic wavelet transform with a flexible time-frequency covering*, IEEE Trans. Signal Process. **61** (2013), no. 5, 1131–1142.
8. Helmut Bölcskei, Franz Hlawatsch, and Hans Feichtinger, *Frame-theoretic analysis of oversampled filter banks*, IEEE Trans. Signal Process. **46** (1998), no. 12, 3256–3268.
9. Z. Cvetković and J. D. Johnston, *Nonuniform oversampled filter banks for audio signal processing*, IEEE Speech Audio Process. **11** (2003), no. 5, 393–399.
10. Zoran Cvetković and Martin Vetterli, *Oversampled filter banks*, IEEE Trans. Signal Process. **46** (1998), no. 5, 1245–1255.
11. Olivier Derrien, Thibaud Necciari, and Peter Balazs, *A quasi-orthogonal, invertible, and perceptually relevant time-frequency transform for audio coding*, Proc. EUSIPCO (Nice, France), IEEE, August 31 – September 4 2015, pp. 804–808.
12. D.L. Donoho, *De-noising by soft-thresholding*, IEEE Trans. Inf. Theory **41** (1995), no. 3, 613–627.
13. Monika Dörfler and Ewa Matusiak, *Nonstationary gabor frames - existence and construction*, IJWMIP **12** (2014), no. 3.
14. Matthew Fickus, Melody L. Massar, and Dustin G. Mixon, *Finite frames and filter banks*, Finite Frames, Applied and Numerical Harmonic Analysis, Birkhäuser Boston, Cambridge, MA, USA, 2013, pp. 337–379.
15. P. Flandrin, *Time-frequency/time-scale analysis*, Wavelet analysis and its application, vol. 10, Academic Press, San Diego, CA, USA, 1999.
16. Bin Gao, W. L. Woo, and L. C. Khor, *Cochleagram-based audio pattern separation using two-dimensional non-negative matrix factorization with automatic sparsity adaptation*, J. Acoust. Soc. Am. **135** (2014), no. 3, 1171–1185.
17. J.S. Garofolo, L.F. Lamel, W.M. Fisher, J.G. Fiscus, D.S. Pallett, and N.L. Dahlgren, *TIMIT acoustic-phonetic continuous speech corpus LDC93S1*, Philadelphia: Linguistic Data Consortium, 1993.
18. B. R. Glasberg and B. C. J. Moore, *Derivation of auditory filter shapes from notched-noise data*, Hear. Res. **47** (1990), 103–138.
19. K. Gröchenig, *Acceleration of the frame algorithm*, IEEE Trans. Signal Process. **41** (1993), no. 12, 3331–3340.
20. Teddy Surya Gunawan, Eliathamby Ambikairajah, and Julien Epps, *Perceptual speech enhancement exploiting temporal masking properties of human auditory system*, Speech Commun. **52** (2010), no. 5, 381 – 393.
21. John HL Hansen and Bryan L Pellom, *An effective quality evaluation protocol for speech enhancement algorithms.*, Proc. ICSLP (Sydney, Australia), vol. 7, November 1998, pp. 2819–2822.
22. V. Hohmann, *Frequency analysis and synthesis using a gammatone filterbank*, Acta Acust. united Ac. **88** (2002), no. 3, 433–442.
23. N. Holighaus, M. Dörfler, G. Velasco, and T. Grill, *A framework for invertible, real-time constant-Q transforms*, IEEE Audio, Speech, Language Process. **21** (2013), no. 4, 775–785.
24. T. Irino and R. D. Patterson, *A dynamic compressive gammachirp auditory filterbank*, IEEE Audio, Speech, Language Process. **14** (2006), no. 6, 2222–2232.
25. Jelena Kovačević and M. Vetterli, *Perfect reconstruction filter banks with rational sampling factors*, IEEE Trans. Signal Process. **41** (1993), no. 6, 2047–2066.
26. J. Le Roux and E. Vincent, *Consistent wiener filtering for audio source separation*, Signal Processing Letters, IEEE **20** (2013), no. 3, 217–220.
27. L. Lin, W.H. Holmes, and E. Ambikairajah, *Auditory filter bank inversion*, Proc. ISCAS (Sydney, Australia), vol. 2, IEEE, May, 6–9 2001, pp. 537–540.
28. E. A. Lopez-Poveda and R. Meddis, *A human nonlinear filterbank*, J. Acoust. Soc. Am. **110** (2001), no. 6, 3107–3118.
29. R.F. Lyon, A.G. Katsiamis, and E.M. Drakakis, *History and future of auditory filter models*, Proc. ISCAS (Paris, France), IEEE, June 2010, pp. 3809–3812.
30. P. Majdak, P. Balazs, and B. Laback, *Multiple exponential sweep method for fast measurement of head related transferfunctions*, J. Audio Eng. Soc. **55** (2007), no. 7/8, 623–637.

31. Piotr Majdak, Peter Balazs, Wolfgang Kreuzer, and Monika Dörfler, *A time-frequency method for increasing the signal-to-noise ratio in system identification with exponential sweeps*, Proc. ICASSP, 2011.
32. D. Marelli and P. Balazs, *On pole-zero model estimation methods minimizing a logarithmic criterion for speech analysis*, IEEE Trans. Audio, Speech, Language Process. **18** (2010), no. 2, 237–248.
33. B. C. J. Moore, *An introduction to the psychology of hearing*, sixth ed., Emerald Group Publishing, Bingley, UK, 2012.
34. Thibaud Necciari, *Auditory time-frequency masking: Psychoacoustical measures and application to the analysis-synthesis of sound signals*, Degree of Doctor of Acoustics, Aix-Marseille University, France, 2010.
35. Thibaud Necciari, Peter Balazs, Nicki Holighaus, and Peter Søndergaard, *The ERBlet transform: An auditory-based time-frequency representation with perfect reconstruction*, Proc. ICASSP (Vancouver, Canada), IEEE, May 2013, pp. 498–502.
36. J. J. O'Donovan and D. J. Furlong, *Perceptually motivated time-frequency analysis*, J. Acoust. Soc. Am. **117** (2005), no. 1, 250–262.
37. Douglas O'shaughnessy, *Speech communication: human and machine*, Addison-Wesley, 1987.
38. Hélène Papadopoulos and Matthieu Kowalski, *Sparse and structured decomposition of audio signals on hybrid dictionaries using musical priors*, J. Acoust. Soc. Am. **134** (2013), no. 1, 666–685.
39. Roy D. Patterson, K. Robinson, J. Holdsworth, D. McKeown, C. Zhang, and Mike H. Allerhand, *Complex sounds and auditory images*, Auditory physiology and perception, Proceedings of the 9th International Symposium on Hearing (Oxford, UK), Pergamond, 1992, pp. 429–446.
40. Zdeněk Průša, Peter L. Søndergaard, Peter Balazs, and Nicki Holighaus, *LTFAT: A Matlab/Octave toolbox for sound processing*, Proc. CMMR (Marseille, France), October 2013, pp. 299–314.
41. Zdeněk Průša, Peter L. Søndergaard, Nicki Holighaus, Christoph Wiesmeyr, and Peter Balazs, *The Large Time-Frequency Analysis Toolbox 2.0*, Sound, Music, and Motion (Mitsuko Aramaki, Olivier Derrien, Richard Kronland-Martinet, and Sølvi Ystad, eds.), Lecture Notes in Computer Science, Springer International Publishing, 2014, pp. 419–442.
42. Seyed Omid Sadjadi and John H.L. Hansen, *Mean hilbert envelope coefficients (mhec) for robust speaker and language identification*, Speech Commun. **72** (2015), 138–148.
43. Christian Schörkhuber, Anssi Klapuri, Nicki Holighaus, and Monika Dörfler, *A Matlab toolbox for efficient perfect reconstruction time-frequency transforms with log-frequency resolution*, Audio Engineering Society Conference: 53rd International Conference: Semantic Audio, AES, January 2014.
44. Elizabeth Shriberg, *Higher-level features in speaker recognition*, Speaker Classification I (Christian Müller, ed.), Lecture Notes in Computer Science, vol. 4343, Springer Berlin Heidelberg, 2007, pp. 241–259.
45. Adrien Sirdey, Olivier Derrien, and Richard Kronland-Martinet, *Adjusting the Spectral Envelope Evolution of Transposed Sounds with Gabor Mask Prototypes*, Proc. DAFx-10 (Graz, Austria), September 2010, pp. 1–7.
46. Diana T. Stoeva and Peter Balazs, *Invertibility of multipliers*, Appl. Comput. Harmon. Anal. **33** (2012), no. 2, 292–299.
47. Stefan Strahl and Alfred Mertins, *Analysis and design of gammatone signal models*, J. Acoust. Soc. Am. **126** (2009), no. 5, 2379–2389.
48. L. N. Trefethen and D. Bau III, *Numerical linear algebra*, SIAM, Philadelphia, PA, USA, 1997.
49. Masashi Unoki, Toshio Irino, Brian Glasberg, Brian C. J. Moore, and Roy D. Patterson, *Comparison of the roex and gammachirp filters as representations of the auditory filter*, J. Acoust. Soc. Am. **120** (2006), no. 3, 1474–1492.
50. P.P. Vaidyanathan, *Multirate systems and filter banks*, Electrical engineering. Electronic and digital design, Prentice Hall, Englewood Cliffs, NJ, USA, 1993.
51. X. Valero and F. Alias, *Gammatone cepstral coefficients: Biologically inspired features for non-speech audio classification*, IEEE Trans. Multimedia **14** (2012), no. 6, 1684–1689.
52. Arun Venkitaraman, Aniruddha Adiga, and Chandra Sekhar Seelamantula, *Auditory-motivated gammatone wavelet transform*, Signal Process. **94** (2014), 608–619.

53. E. Vincent, R. Gribonval, and C. Fevotte, *Performance measurement in blind audio source separation*, IEEE Trans. Audio, Speech, Language Process. **14** (2006), no. 4, 1462–1469.
54. Xiaojia Zhao, Yang Shao, and DeLiang Wang, *Casa-based robust speaker identification*, IEEE Trans. Audio, Speech, Language Process. **20** (2012), no. 5, 1608–1616.
55. Muhammad S. A. Zilany and Ian C. Bruce, *Modeling auditory-nerve responses for high sound pressure levels in the normal and impaired auditory periphery*, J. Acoust. Soc. Am. **120** (2006), no. 3, 1446–1466.
56. Eberhard Zwicker and E. Terhardt, *Analytical expressions for critical-band rate and critical bandwidth as a function of frequency*, J. Acoust. Soc. Am. **68** (1980), no. 5, 1523–1525.

Acoustics Research Institute Austrian Academy of Sciences, Wohllebengasse 12–14, A-1040 Vienna, Austria

E-mail address: {thibaud.necciari,nicki.holighaus,peter.balazs,zdenek.prusa}@oeaw.ac.at